AFRL-SR-AR-TR-10-0022

REPORT DOCUMENTATION PAGE

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ORGANIZATION.

1. REPORT DATE *(DD-MM-YYYY)*	2. REPORT TYPE	3. DATES COVERED *(From - To)*
28-02-2010	Final Technical Report	March 2007 - Feb. 2010

4. TITLE AND SUBTITLE
Cognitive Radar

5a. CONTRACT NUMBER
FA9550-07-1-0182

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)
Goodman, Nathan, A.

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
The University of Arizona
888 N EUCLID AVE
TUCSON AZ 85721-0001

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
Dr. Jon Sjogren
Air Force Office of Scientific Research
875 N. Randolph Street
Suite 325, Room 3112
Arlington, VA 22203-1768

NE

10. SPONSOR/MONITOR'S ACRONYM(S)
AFOSR

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Public Release

20100427016

13. SUPPLEMENTARY NOTES

14. ABSTRACT

Several advances were made toward a foundation for cognitive radar. Several extensions to optimum or matched waveform theory were completed, including formalization of a random-target variance function used in the design methods, extensions to MIMO radar for target identification, information-based waveforms in the presence of ground clutter, incorporation of constant-modulus design techniques, and an adaptive PRF selection technique. These techniques were also applied to spatial waveform design (i.e. beamshaping) in order to develop the fundamentals for a cooperative multiplatform air-to-ground surveillance capability. Two techniques based on the covariance of target track states were developed for integrating detection and tracking into the same Bayesian framework, as well as probability updating techniques in target parameter space for multi-platform detection and tracking. This allowed for beamsteering toward areas in a scene where target presence and/or parameters were most uncertain.

15. SUBJECT TERMS
Cognitive Radar; Radar Detection; Radar Tracking; Radar Surveillance; Ground Moving Target Identification; Target Identification; Waveform Design

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			19b. TELEPHONE NUMBER *(Include area code)*
U	U	U	UU		

Cognitive Radar

Final Performance Report

For

AFOSR Contract FA9550-07-1-0182

Submitted to:
Dr. Jon Sjogren, AFOSR
875 N. Randolph Street
Suite 325, Room 3112
Arlington, VA 22203-1768

Submitted by:
Nathan A. Goodman, Associate Professor
Department of Electrical and Computer Engineering
The University of Arizona

1230 E. Speedway Blvd
Tucson, AZ 85721
Phone: 520-621-4462
Fax: 520-626-3144
goodman@ece.arizona.edu

THE UNIVERSITY OF
ARIZONA
TUCSON ARIZONA

1. GRANT OBJECTIVES

The effort's objectives, which are unchanged from the original SOW, are:

- Develop a Bayesian framework for the implementation of cognitive radar
  - Design an efficient state-vector representation of current information about target tracks, clutter, and external interference
  - Develop an efficient formulation of channel hypotheses that admits the potential for undetected targets, thus integrating detection with other radar functions
  - Quantify relative probabilities of competing channel hypotheses for a complete Bayesian representation of the channel
  - Apply target motion models to generate look-ahead Bayesian channel representations that will allow forward planning of radar positioning
- Extend current work on optimized waveforms for multi-hypothesis testing to support cognitive radar objectives
  - Extend waterfilling-based waveform design technique to account for signal-dependent interferences sources such as ground clutter
  - Develop optimized waveform design strategies for sequential and simultaneous illumination by multiple transmitters
  - Design and test sub-optimum waveform design strategies that do not require a full waterfilling solution (for both signal-dependent and signal-independent models)
  - Design optimum space-time waveforms for discrete clutter cancellation and RF interference reduction
- Integrate task prioritization strategies into the Bayesian hypothesis testing framework
  - Develop approaches for integrating priorities directly into the Bayesian channel representation
  - Develop approaches for integrating priorities into the waveform design
- Demonstrate the newly developed technologies through simulation
  - Develop an air-to-air simulation for demonstration of the Bayesian framework, waveform design strategies, and task prioritization
  - Develop an air-to-ground simulation that further demonstrates the capabilities of a cognitive radar system to perform look-ahead planning and space-time transmit beamforming

- Demonstrate cognitive radar's look-ahead planning capability through an urban GMTI simulation where the radar platforms attempt to retain line-of-sight contact with a target

- When possible, work with AFRL to demonstrate cognitive radar technologies using real data.

2. Executive Summary of Effort

During the "Cognitive Radar" effort, we were successful in making several advances toward a foundation for cognitive radar. First, we made several extensions to optimum or matched waveform theory, including extensions to MIMO radar for target identification, information-based waveforms in the presence of ground clutter, incorporation of constant-modulus design techniques, and an adaptive PRF selection technique. We also applied several of these techniques to spatial waveform design (i.e. beamshaping) in order to develop the fundamentals for a cooperative multiplatform air-to-ground surveillance capability. We developed two techniques for integrating detection and tracking into the same Bayesian framework, as well as probability update techniques in target parameter space for multi-platform detection and tracking. Applying waveform design in clutter techniques to the multi-platform detection and tracking scenario still continues and will be complete soon.

We made significant progress on nearly all objectives including development of the Bayesian framework, extension of waterfilling-based waveform design to account for signal-dependent ground clutter, simultaneous illumination strategies by multiple transmitters, integration of task priorities into the Bayesian framework, and demonstrations through simulation. We also incorporated published literature on waveform design with modulus constraints into our closed-loop framework for target identification.

Personnel who were supported by this project or otherwise worked on this project include:

- Dr. Nathan A. Goodman, Associate Professor and PI, ECE Dept.;
- Mr. Ric Romero, Ph.D. student, ECE Dept., supported by graduate research appointment (GRA);
- Mr. Thomas B. Butler, M.S. student, ECE Dept., supported by GRA;
- Mr. Jun Hyeong-Bae, M.S./Ph.D. student, supported by GRA;
- Mr. Hyoung-soo Kim, Ph.D. student, supported by GRA;
- Mr. Pete Nielsen, M.S. and M.B.A. student, ECE and Business, affiliated with project;
- Mr. Christopher Kenyon, M.S. student, affiliated with project.

At least 12 publications were based on the work supported by this grant. These publications are listed later in the report.

3. ACCOMPLISHMENTS AND NEW FINDINGS

In the following, we first describe a comprehensive theory for waveform design according to SNR and mutual information (MI) metrics. Then we develop a Bayesian framework for target identification, including incorporation of our waveform theory. Finally, we describe our Bayesian approach to surveillance radar, including integrated search and track with multiple platforms.

A. Comprehensive Theory for Matched Waveform Design

The literature has several examples of matched illumination waveform design techniques. Typical design metrics include SNR and MI. We have expanded on these results by developing a comprehensive approach that handles both metrics in a similar manner and can handle signal-dependent clutter. Furthermore, we have derived optimum waveforms for scenarios that were not previously published, included SNR-based waveform design for an ensemble of targets (rather than a single known target), and MI-based waveform design in signal-dependent clutter. Note that our approach to waveform design in signal-dependent clutter does not require an iterative search procedure as does earlier approaches by other authors.

Existing literature in matched waveform design includes matching to a known target for optimized SNR in white noise [1] and in signal-dependent clutter [2], detection of targets in clutter [3], and optimized MI in Gaussian noise [1]. Thus, SNR and MI have been the primary design metrics for matched waveforms, and various target and interference paradigms are possible. We developed waveform design techniques for four new cases: MI-based waveform design in signal-dependent clutter, frequency-domain approach to maximizing SNR in signal-dependent clutter (non-iterative method), and two SNR-based waveform design strategies that apply to target ensembles. Below, we first address the known-target case. We then derive a finite-duration random target model and corresponding SNR- and MI-based waveforms. The complete waveform design theory is in [4], which has been accepted for publication.

1. Waveform Design for Known Targets

Consider the signal model shown in Figure 1. Let $g(t)$ be a known complex baseband target impulse response, $x(t)$ be the baseband transmitted waveform, and $n(t)$ be additive complex Gaussian white noise. Gaussian clutter is denoted by the random impulse response $\mathbf{c}(t)$, and the actual received signal due to clutter is $x(t)*\mathbf{c}(t)$; hence, the measured clutter depends on the interaction of the clutter impulse response with the transmitted waveform. The clutter is

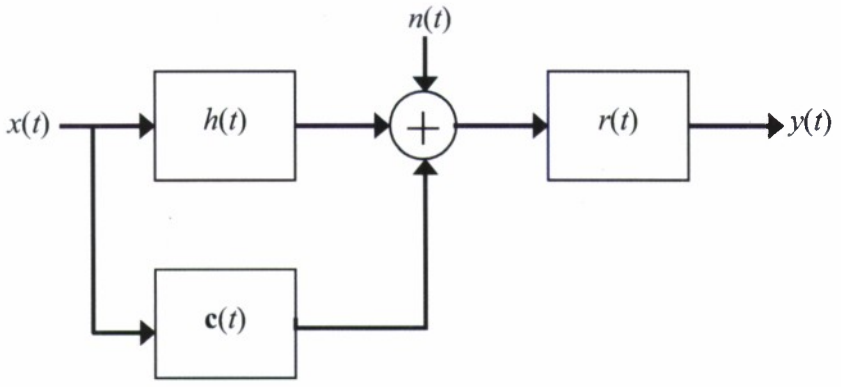

Figure 1. Known-target signal model with clutter and additive white noise.

statistically characterized by a power spectral density (PSD), $S_{cc}(f)$. After the signal, clutter, and noise components are added, the combined signal is passed through the filter with impulse $r(t)$ and sampled at time $t = t_0$. We wish to design the transmit waveform and receive filter that maximize signal-to-interference-plus-noise ratio (SINR) at the sample time.

Let the transmit waveform's energy be constrained such that .

$$E_x = \int_{-\infty}^{\infty} |x(t)|^2 \, dt = \int_{-\infty}^{\infty} |X(f)|^2 \, df \, . \tag{1}$$

The SINR can be expressed as

$$\mathrm{SINR}_{t_0} = \frac{\left| \int_{-\infty}^{\infty} R(f) H(f) X(f) e^{j2\pi f t_0} df \right|^2}{\int_{-\infty}^{\infty} |R(f)|^2 L(f) df} \tag{2}$$

where $L(f) = |X(f)|^2 S_{cc}(f) + S_{nn}(f)$. Using Schwarz's inequality and then the method of Lagrange multipliers, the optimum waveform spectrum is

$$|X(f)|^2 = \max\left[0, B(f)\left(A - D(f)\right)\right] \tag{3}$$

where

$$B(f) = \frac{\sqrt{|H(f)|^2 S_{nn}(f)}}{S_{cc}(f)}, \tag{4}$$

$$D(f) = \frac{\sqrt{S_{nn}(f)}}{\sqrt{|H(f)|^2}}, \tag{5}$$

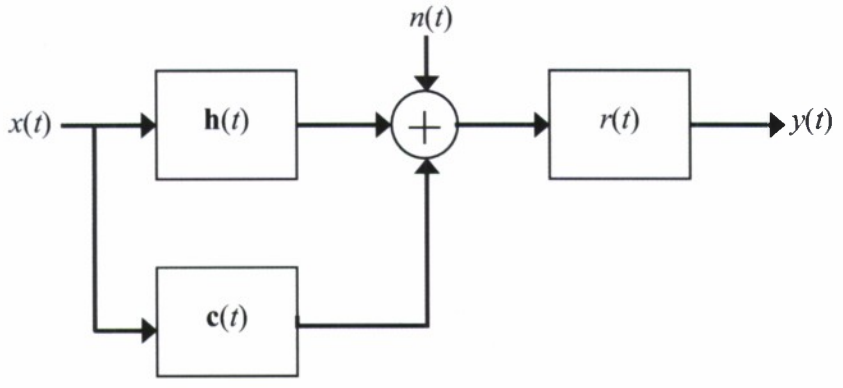

Figure 2. Random-target signal model with clutter and additive white noise.

and A is a constant that is determined according to the energy constraint. The complete derivation can be seen in [4].

The solution (3) is found through the numerical procedure known as waterfilling, which is much easier to solve than earlier iterative solutions to this problem. Furthermore, if clutter is set to zero (i.e., $S_{cc}(f) = 0$), then the expression for SINR converges to the additive noise case presented in [1] where the waveform is found by finding the primary eigenvector of an integral equation .

2. Finite-Duration Random Target Model

Bell [1] derived an information-based waveform using an extended target model where the random target impulse response $\mathbf{h}(t)$ was stated to be a finite-energy, finite-duration random process with something called a spectral variance. The meaning of a finite-duration random process was not formalized in [1]; however, we have formalized the definition as part of our waveform design work performed under this grant. Figure 2 shows the signal model for a random impulse response in the presence of additive noise and signal-dependent clutter.

Let $\mathbf{g}(t)$ be a true wide-sense stationary complex random process (because it is stationary, it must be infinite duration) with some PSD and let $a(t)$ be a rectangular window of duration T_h. The product $\mathbf{h}(t) = a(t)\mathbf{g}(t)$ is a finite-duration random process having support only in $[0, T_h]$. The statistics for $\mathbf{h}(t)$ are stationary within its interval of support. Since $\mathbf{h}(t)$ has finite support, it is a finite-energy process, and any realization $h(t)$ will be integrable. As such, we can define a Fourier transform $H(f)$ and state the energy relationship

$$E_h = \int_{T_h} |h(t)|^2 \, dt = \int_{-\infty}^{\infty} |H(f)|^2 \, df \, . \tag{6}$$

From (6) and the known duration of the target, we can define a time-averaged power P_h that is valid within the random target's support interval. Applying this time average and taking the expected value with respect to the target ensemble, we have

$$\overline{P}_{\mathbf{h}} = \frac{1}{T_h}\int_{T_h} \mathrm{E}\left[\left|\mathbf{h}(t)\right|^2\right] dt = \frac{1}{T_h}\int_{-\infty}^{\infty} \mathrm{E}\left[\left|\mathbf{H}(f)\right|^2\right] df \,. \tag{7}$$

We define $\mathrm{E}\left[\left|\mathbf{H}(f)\right|^2\right]$ as an energy spectral density (ESD) for the finite-duration random target and

$$\sigma_H^2(f) = \mathrm{E}\left[\left|\mathbf{H}(f) - \mathrm{E}\left[\mathbf{H}(f)\right]\right|^2\right] \tag{8}$$

as the energy spectral variance (ESV). With the time averaging, we have (assuming zero mean)

$$\Upsilon_H(f) = \frac{\sigma_H^2(f)}{T_h} = \frac{\mathrm{E}\left[\left|\mathbf{H}(f)\right|^2\right]}{T_h}, \tag{9}$$

which we call the power spectral variance (PSV). Note that $\Upsilon_H(f)$ is not a power spectral density because $\mathbf{h}(t)$ is not a true power signal or stationary random process. However, $\Upsilon_H(f)$ conveys the same basic information as a PSD, except that it is only valid within the finite support interval.

Next, we must determine what happens when this finite-duration random target is convolved with a finite-duration waveform. We define the output of the convolution as

$$\mathbf{z}(t) = x(t) * \mathbf{h}(t). \tag{10}$$

The random output $\mathbf{z}(t)$ only has support on the interval $[0, T_z]$ where $T_z = T + T_h$ and T is the waveform duration; furthermore, $\mathbf{z}(t)$ it is not stationary within that interval due to ramp-up and ramp-down periods of the convolution. This complicates the statistics of the output. We take the approach of defining an ESV for the output signal according to

$$\sigma_Z^2(f) = \left|X(f)\right|^2 \sigma_H^2(f). \tag{11}$$

The time-averaged PSV is then

$$\Upsilon_Z(f) = \frac{\sigma_Z^2(f)}{T_z} = \frac{\left|X(f)\right|^2 \sigma_H^2(f)}{T_z} = \alpha\left|X(f)\right|^2 \Upsilon_H(f). \tag{12}$$

Note that if we let the target become a true random process by allowing $T_h \to \infty$, then the PSV of $\mathbf{h}(t)$ becomes a true PSD, and (12) becomes the input-output relationship that we are accustomed to for stationary random processes passing through a linear, time-invariant filter.

3. SNR-based Waveform Design for Random Targets

We use the local SNR metric defined in [5] for a true random process $\mathbf{g}(t)$, then substitute the finite-duration relationships. Allowing the target to be infinite duration, the output SINR spectral density is

$$R_{SINR} = \frac{S_{gg}(f)\left|X(f)\right|^2}{\left|X(f)\right|^2 S_{cc}(f) + S_{nn}(f)}, \tag{13}$$

and the SINR is achieved by integrating (13) over all frequencies. If the signal is observed and integrated over an interval T_0, then

$$\text{SINR} = T_0 \int_{-\infty}^{\infty} R_{SINR}(f)\,df\,. \tag{14}$$

Substituting our time-averaged quantities from the finite-duration target model, the SINR is

$$\text{SINR} = \int_{-\infty}^{\infty} \frac{\sigma_H^2(f)\left|X(f)\right|^2}{\left|X(f)\right|^2 S_{cc}(f) + S_{nn}(f)}\,df\,. \tag{15}$$

Applying the Lagrange multiplier technique, the waveform spectrum that maximizes (15) is

$$\left|X(f)\right|^2 = \max\left[0, B(f)\left(A - D(f)\right)\right] \tag{16}$$

where

$$B(f) = \frac{\sqrt{\sigma_H^2(f) S_{nn}(f)}}{S_{cc}(f)}, \tag{17}$$

$$D(f) = \sqrt{\frac{S_{nn}(f)}{\sigma_H^2(f)}}, \tag{18}$$

and A is determined from the finite-energy constraint.

4. MI-based Waveform Design for Random Targets

In this design approach, we wish to maximize the mutual information between the output signal $\mathbf{z}(t)$ and the random impulse response $\mathbf{h}(t)$. First, using the infinite-duration target $\mathbf{g}(t)$, the mutual information between output signal and target would be infinite because $\mathbf{g}(t)$ has infinite

entropy owing to its infinite duration (consider the number of time samples and bits necessary to represent a realization of $\mathbf{g}(t)$). What can be defined then, is an information rate, which is

$$\dot{I}\left(\mathbf{z}(t);\mathbf{g}(t)\,|\,x(t)\right)=\int_W \ln\left[1+\frac{|X(f)|^2 S_{gg}(f)}{|X(f)|^2 S_{cc}(f)+S_{nn}(f)}\right]df. \tag{19}$$

Using our finite-duration substitutions and multiplying by the output interval T_z, an approximate expression for the finite-duration case is

$$I\left(\mathbf{z}(t);\mathbf{h}(t)\,|\,x(t)\right)=T_z\int_W \ln\left[1+\alpha\frac{|X(f)|^2 \Upsilon_H(f)}{|X(f)|^2 S_{cc}(f)+S_{nn}(f)}\right]df. \tag{20}$$

where $\alpha=T_h/T_z$. Insofar as we are willing to take the analogy between a true PSD and the time-averaged PSV, the expression in (20) is equivalent to (19) (plus the output duration term T_z in front). Optimizing (20) via Lagrange multipliers gives the waveform spectrum

$$|X(f)|^2=\max\left[0,-R(f)+\sqrt{R^2(f)+S(f)\left(A-D(f)\right)}\right] \tag{21}$$

where

$$D(f)=\frac{S_{nn}(f)}{\alpha\Upsilon_H(f)} \tag{22}$$

$$R(f)=\frac{S_{nn}(f)\left(2S_{cc}(f)+\alpha\Upsilon_H(f)\right)}{2S_{cc}(f)\left(S_{cc}(f)+\alpha\Upsilon_H(f)\right)} \tag{23}$$

$$S(f)=\frac{S_{nn}(f)\alpha\Upsilon_H(f)}{S_{cc}(f)\left(S_{cc}(f)+\alpha\Upsilon_H(f)\right)}, \tag{24}$$

and is determined from the waveform energy constraint.

The waveform theory has a nice symmetry between the equations for the different cases. We have strengthened the definitions of spectral variance used in [1], and if the clutter PSD is set to zero in the MI-based waveform design, the equations reduce to the same equations used in [1]. Thus, the theory behaves as expected with respect to other published literature and as the duration of the finite target is allowed to go to infinity. Examples of waveform spectra and further commentary on the relationships between the cases can be found in [4,6].

These equations provide nice solutions to optimum waveform spectra under two different design metrics and for varying interference scenarios. The next step is to determine how the

customized waveforms can be used effectively in learning about a probabilistically rated radar channel.

B. Closed-Loop/Cognitive Radar Framework for Radar Target Identification

We have developed a closed-loop radar interrogation framework that can be applied to target identification applications, as well as integrated detection and target tracking. This closed-loop framework integrates matched illumination waveform design techniques with a Bayesian representation of the radar channel. The Bayesian representation characterizes the fundamental and relevant uncertainty that the radar system has about its propagation environment. For example, if the goal of the radar system is to classify a target, then the Bayesian model reduces to an ensemble of target possibilities for each target class as well as a pdf describing the target realizations in each class. The Bayesian representation can then be combined with the waveform design techniques above to transmit a waveform that is fine tuned to the fundamental question of target classification. In other words, the Bayesian model represents the radar channel as a probabilistically rated set of hypotheses, and once hypotheses have been defined, waveforms can be designed to enhance decision-making performance. Moreover, the interrogation procedure is adaptive rather than fixed. While the Bayesian representation can be used to fine tune waveforms, the received observations can be used to update the Bayesian representation. Hence, future interrogations depend on past and present received data via updates of the Bayesian representation.

The UA has demonstrated that a closed-loop framework with adaptive waveform design can significantly decrease the number of transmissions necessary for performing statistical target recognition with some specified accuracy, or equivalently, can improve accuracy for a given number of transmissions. The closed-loop framework currently represents different hypotheses as either known impulse responses or as target classes, depending on the application and prior knowledge. In both cases, a prior probability is assigned to each hypothesis, and in the case of target classes, each class has a pdf that describes the relative likelihood of different realizations in the class. Hence, the probabilistic model inherently incorporates prior knowledge through class pdf's and prior probabilities. In the following, we describe the UA closed-loop radar framework in more detail.

1. Problem Setup and Signal Model

Consider the hypothesis testing problem that was evaluated in [7]. Let there be M known hypotheses. Each hypothesis represents a different linear system, such as a different channel or target. Let the mth hypothesis be characterized by an impulse response $h_m(t)$ for $m \in \{1,2,\ldots,M\}$. Therefore, if the system transmits a waveform $x(t)$, then the received signal under the m^{th} hypothesis is

$$y(t) = x(t) * h_m(t) + n(t) \tag{25}$$

where $*$ is the convolution operator and $n(t)$ is additive noise. The decision to be made is to choose the correct impulse response of the system. The cognitive radar system computes a waveform that is matched to the ensemble of hypotheses with initial probabilities. The waveform is transmitted, and the observations are made. The system then computes likelihood ratios between the various hypotheses and compares them to thresholds determined from the theory of sequential hypothesis testing. If the likelihood ratios corresponding to one of the hypotheses all exceed their predetermined thresholds, then a decision is made for that hypothesis. If no decision can be made, then the hypothesis probabilities are updated using Bayes' rule. Finally, the updated probabilities are used to compute a waveform that is matched to the ensemble with new probabilistic ratings, the waveform is transmitted, and the process continuous until a decision is made.

When combined, the hypothesis probability updates and the matched waveform design strategy form the basis for a closed-loop radar system that continually adapts its waveform to improve discrimination between the most likely remaining hypotheses. This closed-loop process continues until a sequential hypothesis testing strategy determines that sufficient confidence for a decision has been achieved. We apply this closed-loop framework to scenarios where the hypotheses are known impulse responses, to scenarios where the hypotheses are statistically characterized classes of impulse responses, and to scenarios where targets are described by a library of impulses responses that describe the target from different aspect angles.

2. Application of Waveform Designs

Now that a target identification problem has been defined, we need to determine how the waveform design equations derived above might aid in identifying a target more quickly or with reduced time and/or energy. The key is to recognize that a Bayesian ensemble of potential targets

can be converted into a variance function that describes the variability of the different targets as a function of frequency. Once the variance function is defined, it can be substituted into one of the waveform design equations above.

a) Ensemble of Known Target Impulse Responses

Suppose that the target impulse responses, the $h_m(t)$'s, are known and that the probability that the mth target is the true target is P_m. In this case, a spectral variance is straightforward to define [7]:

$$\sigma_H^2(f) = \sum_{m=1}^{M} P_m \left|H_m(f)\right|^2 - \left|\sum_{m=1}^{M} P_m H_m(f)\right|^2 . \tag{26}$$

This spectral variance can then be substituted into one of the waveform design techniques above in order to spectrally shape a waveform with good discriminating ability. Both the SNR-based and MI-based waveforms will focus their finite energy on the frequencies where the spectral variance from (26) is high and, therefore, where the targets are most identifiable. As the probabilities are updated in response to received measurements, the spectral variance in (26) obviously changes, resulting in a newly optimized waveform.

Bayes' rule can be used to update the probabilities in response to received measurements. Suppose that the received signal is $z(t)$, then Bayes' rule says that

$$P_m \mid z(t) = \frac{p_m\left(z(t)\right)P_m}{p\left(z(t)\right)} \tag{27}$$

where $p\left(z(t)\right)$ is the pdf of the measurements and $p_m\left(z(t)\right)$ is the pdf of the measurements conditioned on the mth hypothesis. The denominator term in (27) is common to all hypotheses and serves only to scale the probabilities such that they sum to unity. For the case in this section, the target impulse response/transfer function is known for each hypothesis, so the received signal under the mth hypothesis has a mean value given by $z(t) = x(t) * h_m(t)$. The full pdf is determined by the statistics of the additive noise and clutter.

b) Ensemble of Random Targets

In most cases, the target hypotheses will not be known exactly, but rather each hypothesis is for a particular target type with unknown orientation. Thus it would be more reasonable to describe each hypothesis statistically with a pdf for its impulse response or transfer function. Let

$\sigma^2_{H,m}(f)$ be the spectral variance that defines the mth target class hypothesis. The equation that seems to be consistent with (26) is

$$\sigma^2_H(f) = \sum_{m=1}^{M} P_m \sigma^2_{H,m}(f) - \left| \sum_{m=1}^{M} P_m \sqrt{\sigma^2_{H,m}(f)} \right|^2, \tag{28}$$

which may be thought of as an effective ESV over multiple target classes. We have obtained good results by substituting (28) into our waveform design equations, but the result is not altogether satisfying due to the unclear meaning of the root of the spectral variance. Thus, we have also tried an approach where we design waveforms that are matched to each target class, then weight the energy in those waveforms according to P_m. This strategy was reported in [8].

Bayes' rule can again be used in this scenario, but the pdf of the measured data as well as the conditional pdf of the measured data differs due to the random nature of the target. Furthermore, different pdfs can result from multiple transmissions/receptions depending on whether the target realization is assumed to stay constant over the transmissions or whether it fluctuates over the transmissions. Probability density functions for several different scenarios of known/random target; correlated or independent over multiple transmissions; and correlated or independent clutter realizations have been described for the Gaussian case in [6] and need not be re-derived here.

3. Results

We have tested the above waveform derivations and target variance models using the closed-loop target recognition approach. In this approach, the radar system assigns initial probabilities to several possible target hypotheses. Each target class is characterized either by a different known impulse response or by a PSV. The true target hypothesis is randomly selected, and in the random-target case, a target realization from the selected class is randomly generated. The radar system transmits a waveform that is matched to the target hypotheses according to the hypothesis probabilities, collects the observations, and uses sequential hypothesis testing theory to determine if a decision can be made. If not, the probability of each hypothesis is updated, a new waveform is computed, and the process continues. The number of transmissions needed to make a decision is a random variable since different noise and target realizations affect when a decision can be made.

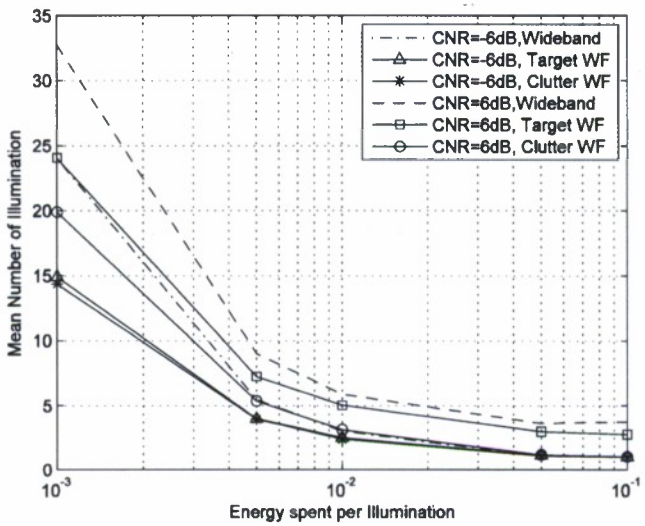

Figure 3. Average number of transmissions to make a target recognition decision for a target embedded in clutter. Each hypothesis was a class of targets defined by a PSD.

Figure 3 shows a result from one of our simulations. The adaptive interrogation procedure was repeated 1500 times. For each trial, one of the target classes was randomly selected as the true hypothesis, and a realization of the true target class was generated. A realization of the ground clutter impulse response was also generated. The closed-loop procedure was performed until a decision was made, and for each trial, the number of transmissions was noted. Then, the average number of iterations over the 1500 trials was computed. We see in Figure 3 that the MI-based waveform strategy that accounts for the ground clutter PSD ("Clutter WF") outperforms the other waveform strategies, especially for higher clutter-to-noise ratio (CNR). The curve labeled "Target WF" is obtained through MI-based waveform design that is matched to the target ensemble's spectral variance and adapting the waveform at each transmission, but not accounting for the clutter spectrum. The "Wideband" waveform is an impulse-like waveform with evenly allocated energy in the frequency domain. Therefore, it is not matched to the target ensemble and does not adapt as the hypothesis probabilities are updated. Additional results are available in [4,6-10].

4. Additional Scenarios and Results

a) MIMO Adaptive Waveform Design for Target Recognition

We have also considered the question of how a multiple-input, multiple-output (MIMO) radar waveform should be designed for the goal of target recognition. The MIMO waveform design problem has been considered in [11]. In [11], the solution is in the time-domain, but the

matrix solution must have a Toeplitz form in order to represent a physical waveform that convolves with the target impulse response to produce a receive echo. This constraint is mentioned, but apparently is never enforced. Hence, we have been unable to implement the solution in [11] in our simulations. In [12], the MIMO context is a bit different – the goal is to transmit different temporal waveforms over different spatial beams in order to improve tracking and parameter estimation for multiple targets. We consider the *orthogonal* transmit waveform approach. In the orthogonal approach, it is assumed that each receiver can separate the received signals due to each of the transmitters. This can be done, for example, if each transmitter uses a non-overlapping frequency band.

Complete details of the MIMO waveform scenario can be found in [13] and in the M.S. Thesis by Thomas Butler, which will be included as an attachment to this report.

b) Constant Modulus Constraints

In our simulation work, a time-domain waveform can be obtained by simply taking the root of the square of the waveform's magnitude spectrum as defined in the waveform design equations, then computing an inverse Fourier transform. However, this approach yields waveforms that are not generally constant modulus in the time domain and, therefore, do not operate continuously at full peak power. Fortunately, the waveform design equations only specify the waveform's magnitude spectrum, and not its phase characteristic. This remaining flexibility can be used to design constant modulus waveforms with magnitude spectra that nearly match the optimum desired spectra. We have not developed these constant-modulus design approaches, be we implemented them into our closed-loop target identification framework and published the results in [14]. Our results show significant performance improvement can still be maintained with constant-modulus waveforms even though the spectrum of the constant-modulus waveform may not exactly matched the optimum spectrum from one of the design equations. Furthermore, if a peak-power constraint is applied to a non-constant-modulus waveform obtained via SNR- or MI-based design, the benefit of an optimized waveform may be reduced by a large amount.

C. Spatial-Domain Matched Illumination for Integrated Search and Track

Traditional radar systems perform search and tracking functions in a very rigid manner. Periodic beams sweep across the search area, and track updates are scheduled at regular intervals. In contrast to this, an adaptive-transmit radar system would use information it has

gathered and stored to dynamically control the search and track processes. For search, the radar system would be able to steer the antenna beam to locations in the propagation environment that had greater uncertainty of target presence or absence. In other words, the radar system would not be forced by the static sweep pattern to search areas where there is a high level of certainty regarding target presence. This would allow for a better allocation of time and energy across the propagation environment. For track, the adaptive radar system would be able to change the update rate based upon the current knowledge about the estimated target parameters. This would allow the radar system to update estimated target parameters more often for fast moving targets and less often for easy-to-track slow moving targets.

We have developed a spatial-domain application of the mutual-information-based matched waveform technique in order to develop and test a closed-loop implementation of target detection and tracking. Through the equivalence of angle of arrival with spatial frequency, we developed a probabilistic representation of potential target locations in target parameter space. This probabilistic representation can then be converted into a two-dimensional *spectral variance function* (a function of spatial frequencies k_x and k_y) upon which the waterfilling operation can be applied to find a matched transmit beampattern. The target parameter model includes target velocity; therefore, potential targets can move between the time when the ensemble is updated with a data collection and the time in the future when the next transmission will occur. Thus, we also use a Kalman-based prediction step to anticipate the status of the probabilistic channel representation at the time when the transmission will occur. The result is a system with two different techniques for performing fully integrated search and track functions. The system scans and shapes its transmit beam not according to a pre-defined or fixed timeline, but according to the uncertainty of the probabilistic model.

1. Single-Platform Model and Closed-Loop Beam Control

To begin, we employ a simplified single-platform signal model that ignores range resolution and ground clutter. This allows us to focus on the probability-updating adaptive-beamsteering aspects of the system. In this scenario, the propagation channel is defined only by the presence and/or absence of targets, and by the parameters of targets that are present. Right away this idea of mixing detection (target presence/absence) with target tracking (parameter estimation/ uncertainty) poses an interesting problem. The measures of uncertainty in these two cases are completely different – a discrete random variable for the number of targets present and

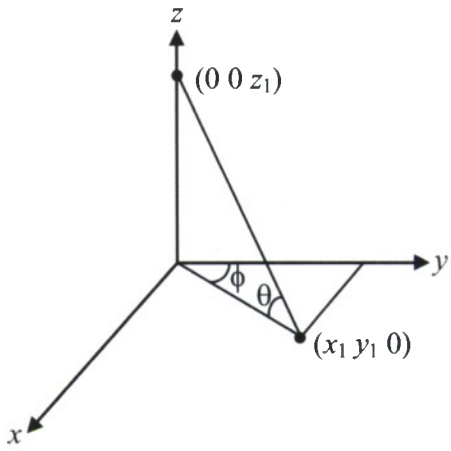

Figure 4. Simplified single-platform look-down geometry.

continuous random variables for the target parameters. A key challenge was how to integrate these two different types of objectives into a single framework for updating the Bayesian channel ensemble and controlling the transmit beam. In effect, there are infinitely many hypotheses for the channel due to the unknown number of targets and their continuum of possible parameters (position, velocity, ...). Toward this end, we propose two different techniques, each with its own advantages and disadvantages.

Consider a "look-down" geometry [15-16] depicted in Figure 4 where targets are restricted to the x-y plane and the radar (assumed stationary) is located at height z_1 on the z-axis. The angles θ and ϕ are the elevation and azimuth angles, respectively, from the radar to the target; therefore, the unit vector pointing from the radar to target is

$$\hat{u}_R = \cos\theta\sin\phi\hat{u}_x + \cos\theta\cos\phi\hat{u}_x - \sin\theta\hat{u}_z . \tag{29}$$

The target velocity, which is restricted to be in the x-y plane, is defined as $\mathbf{v} = v_x\hat{u}_x + v_y\hat{u}_y$. Therefore, the Doppler shift due to the motion between the target and radar platform is

$$F_D = -\frac{2\mathbf{v}\cdot\hat{u}_R}{\lambda} = \frac{-2\left(v_x\cos\theta\sin\phi + v_y\cos\theta\cos\phi\right)}{\lambda}. \tag{30}$$

Finally, we assume that the antenna is able to steer a small beam over a region of the x-y plane located in the radar's far field, but that the region is near to the y-axis such that ϕ is small. In this case, the Doppler shift is approximated as

$$F_D \approx \frac{-2v_y\cos\theta\cos\phi}{\lambda}. \tag{31}$$

The radar antenna is assumed to be a multi-channel system such that angles can be measured. Thus, the radar system is capable of measuring two angles and a Doppler shift. Accordingly, the target parameter space is three-dimensional and we will denote the parameters as k_x, k_y, and F_D. Furthermore, the specific orientation of the antenna array and steering angle determine the relationships between this parameters, which will be used in the system model of the Kalman filter. For example, the Doppler shift equations above imply a relationship between Doppler shift and angle rate.

Now, in order to set up a Bayesian channel model, we must determine how to handle continuous target parameters. Our approach is to *discretize* the target parameter space. Let the span of possible values for k_x, k_y, and F_D be divided into cells that are nominally related to the radar's Rayleigh resolution in those three parameters. This step forms a discrete, three-dimensional target parameter space, and the channel ensemble is formed by considering the potential for target presence or absence in each cell. In other words, when every possible permutation of target presence/absence is considered, this forms the channel ensemble. To complete the probabilistic representation, each cell is assigned a probability of target presence. Each time data are collected, these probabilities will be updated, and when probabilities reach a pre-specific high level, we will begin an integrated tracking function.

As we stated earlier, our waveform design results are useful when a probabilistic channel representation can be converted to a variance function that can be inserted into the waveform design equations. Therefore, now that we have a set of possible channel realizations and a probability associated with each target parameter cell, we must convert the realizations and probabilities into a variance function. To do this, we denote the target reflection coefficient as α and compute the variance of the reflection coefficient across the two hypotheses of the detection problem. Noting that the target reflection coefficient is zero under the target-absent hypothesis, the variance for a single parameter cell is

$$\sigma^2 = P|\alpha|^2(1-P). \tag{32}$$

Clearly, as the probability of target presence goes to either zero or one, the variance goes to zero. This is desirable because a probability of either zero or one denotes certainty about target presence; and when variance goes to zero, the waveform design methods will not allocate any energy to estimating that component.

Now to relate this to our three-dimensional target parameter scenario, we note that two of the target parameters are angles. The other parameters is Doppler shift, but an antenna beam will not have variation over Doppler shift. In other words, a target's illumination is only controlled by its location, not by its Doppler shift. Thus, we must convert the three-dimensional parameter space into a variance function that depends on the two beamsteering angles. To accomplish this, let the ith Doppler bin for a particular k_x-k_y cell have probability $P_i(k_x, k_y)$ and reflection coefficient $\alpha_i\left(k_x, k_y\right)$. Assuming that the hypotheses in different cells are independent (target presence in one cell is not related to target presence in another cell), we can sum the variances over all Doppler bins to achieve the variance function

$$\sigma^2\left(k_x, k_y\right) = \sum_{i=1}^{N} \sigma_i^2\left(k_x, k_y\right) = \sum_{i=1}^{N} P_i\left(k_x, k_y\right)\left|\alpha_i\left(k_x, k_y\right)\right|^2\left(1 - P_i\left(k_x, k_y\right)\right). \tag{33}$$

Other than slight modifications to account for two dimensions, the variance in (33) can be directly inserted into the waveform design equations to determine the approximate mutual information that can be expected from a given beam position. Of course, the waveform design equations could be used to obtain an arbitrary k_x-k_y illumination pattern, but this would be difficult to implement. Instead, we use the MI equations to compute the MI that would be achieved by a particular beam shape as a function of the steering angles. The steering angles that suggest the most benefit in terms of MI are the angles used for the next transmission.

Next, once data are collected, we apply Bayes' rule to update the probabilities in each target parameter cell. The exact pdf for Bayes' rule depends on the target and noise models assumed, but if the cell spacing is related to the underlying radar resolution, then the update largely depends on matched filter output. Thus, in some cases it is possible to use traditional matched filter processing to update the channel ensemble – the difference is that instead of making *one-and-done* decisions regarding target presence, the matched filter output is used in conjunction with prior knowledge about target presence in a given cell. Furthermore, the result of the probability update becomes the prior probability used for the next illumination.

Finally, the above discussions focus on the detection problem, but it was an objective of this project to develop strategies for integration of detection and tracking into a single mode. From the abstract notion of Bayesian channel representations and probabilistically rated hypotheses about targets and their parameters, this objective is intuitively straightforward. Implementation,

however, makes it clear why these functions have always been separated in conventional radars. We have proposed two different techniques for integrating the two radar functions.

Theoretically, the ideal solution would be to propagate the entire Bayesian ensemble forward in time for the next illumination. If there was a 10% chance of a target at a particular location and Doppler shift, then there should be a nearly 10% chance that the same target would be at a new location on the subsequent transmission. Unfortunately, this solution is not easy to visualize its implementation, and any implementation would probably be computationally impractical. Instead, we define a soft detection threshold above which we begin a traditional Kalman tracker. The tracker does not necessarily track targets, but instead tracks cells of high probability in target parameter space, and is used to propagate those cells forward in time to the next transmission.

When a target parameter cell exceeds our soft detection threshold, we initiate a track with a parameter state vector and associated covariance. The covariance depends on SNR and radar resolution. A linear target motion model is used to propagate the state vector and its covariance forward in time. Then, we assume that the state covariance defines an ellipse of uncertainty in target parameter space. This ellipse describes the shape of a pdf that describes the target parameters; therefore, to obtain the probability that the (soft) target is within a particular cell, we integrate the pdf, which is defined by the covariance ellipse, over the boundaries of the cell. When the covariance begins to spread over multiple cells, the probability of target presence gets spread over those cells as well, uncertainty increases, and the beam will be steered in that direction to update the cells and reduce uncertainty.

The major benefit of this approach for integrated detection and tracking is that the integration is smooth and natural. Waveforms (in this case steering a spatial waveform) depend on the probability ensemble, and detection and tracking both reduce to probability of target presence in a given cell. There are two main disadvantages. The first is that by completely integrating the two functions, it is difficult to change the emphasis on either mode. Of course, this shouldn't be surprising since it was our goal to treat them as one objective, not two objectives.

The second disadvantage occurs when the target state vector falls on the boundary of two or more cells. Even if the track covariance is extremely tight, if the estimated parameter is exactly on the boundary, roughly half of the target probability will fall into each cell. Since 50% is as uncertain as we can get, the system will update the target even though the target parameters are known very well. The target will still be on a boundary, however, so the system will continue to

update the target until it moves off the cell boundary. This second disadvantage is a direct consequence of discretizing the parameter space. Variable cell sizes and positions have been considered to handle this problem, but the solution begins to become much more clumsy and loses some of its elegance.

A second approach to integrated detection and tracking is to form a hybrid variance function that consists of uncertainty due to probabilities of target presence in a cell combined with the entropy of the target state vector. The problem here is that the two measures of uncertainty are completely different – the detection variance depends on a binary hypothesis problem in each cell while the tracking variance depends on the variance of continuous target parameters. A scaling factor would be necessary to combine the two variances. On the other hand, this scaling parameter could be the method used to control relative priorities. Depending on the value of the scale factor, the system could emphasize detection of new targets, tracking of existing targets, or a balance between the two.

Early results for our work on single-platform cognitive radar for integrated search and track have been reported in [15]. Complete details of the methods used and additional results are found in the thesis by Nielsen [16], which will be included as an attachment to this report.

2. Two-Platform Model and Probability Updating

We now consider two radar platforms jointly performing search and track over a wide area, but observing the scene from different aspects. This paradigm has the potential for allowing multiple radar platforms to cooperate in ways such as handing off targets as they change directions and pass through clutter ridges. Conceptually, the approach is similar to the single-platform scenario described above, but with a few key differences. First, in the single-platform case, only radial velocity can be measured, so the parameter of interest was Doppler shift. When multiple platforms observe a scene from different angles, the Doppler shift observed by the two platforms is different; hence, the more natural target parameters are the absolute velocity components. Each platform must map velocity hypotheses into the Doppler shift that should be observed from its angle. Second, but related to the first, is that different radars will be able to measure or resolve certain target parameters, but will be blind to other parameters. For example, consider two platforms observing a scene from aspect angles that differ by 90 degrees. One radar may be able to measure v_x, but not v_y while the other can measure v_y, but not v_x. This last point

complicates the probability updating procedure, which is where we focus the following discussion.

But before we address the multi-platform scenario, we must consider more carefully how a single radar measuring a parameter θ with M cells would update its cell probabilities with received measurements. We do not assume that the cells are resolution cells at this time, but instead consider a general case. Let the prior probability for the mth cell be given by $P_{m,0}$. Then, consider a sensor that produces an N-clement measurement vector with each data collection. These measurements might be N slow-time measurements for the purpose of measuring Doppler, N spatial measurements for measuring angle, or any other combination of measurements. Let $\mathbf{s}_m$ be the sampled signal produced at the radar if a target is present in the mth parameter cell, and let the measurement indices be denoted by $n = 0, 1, ..., N\text{-}1$. For a radar system, target paramcters can be described as frequencies (e.g., Doppler, spatial, or range), so we let the frequency produced by the presence of a target in the mth cell be denoted by F_m. When these cell frequencies are the same, the cells cannot be resolved and are considered ambiguous. When the frequencies differ, it may be possible to resolve thc cells depending on the difference between frequencies and the size N of the measurement vector. The signal produced by a target in the mth cell is proportional to a normalized steering vector given by

$$\mathbf{s}_m = \frac{1}{\sqrt{N}} \exp\left(j2\pi F_m \begin{bmatrix} 0 & 1 & \cdots & N-1 \end{bmatrix}^{\mathrm{T}} \right) \tag{34}$$

To perform a probability update, it is necessary to assume a probability model for the measurements. For convenience, let thc targets be dcterministic, i.c., known amplitude and phase and let the noise be additive white Gaussian. Recalling that there arc 2^M possible permutations of the overall target environment, the environment can be described by a multiple-hypothesis framework. The hypotheses in this framework are given by

$$\begin{aligned}
H_0 &: \ \mathbf{z} = \mathbf{n} \\
H_1 &: \ \mathbf{z} = \mathbf{s}_1 + \mathbf{n} \\
H_2 &: \ \mathbf{z} = \mathbf{s}_2 + \mathbf{n} \\
H_3 &: \ \mathbf{z} = \mathbf{s}_1 + \mathbf{s}_2 + \mathbf{n} \\
H_4 &: \ \mathbf{z} = \mathbf{s}_3 + \mathbf{n} \\
&\vdots \\
H_{2^M-1} &: \ \mathbf{z} = \mathbf{s}_1 + \mathbf{s}_2 + \cdots + \mathbf{s}_M + \mathbf{n}
\end{aligned} \tag{35}$$

Each hypothesis H_i may be thought of as a *joint* hypothesis corresponding to a unique permutation of target presence/absence in the individual cells. For notational convenience, we convert the joint hypothesis subscript i to its binary representation of $i = 0\cdots00, 0\cdots01, 0\cdots10, 0\cdots11, \ldots, 1\cdots11$, where a 0 corresponds to target absent and a 1 corresponds to target present in a cell. For example, consider the eleventh hypothesis in a five-cell scenario. If we let S_i correspond to the target signal produced by the ith joint hypothesis, then the eleventh hypothesis would be represented as H_{01011} and the received signal contribution would be $S_{01011} = \mathbf{s}_4 + \mathbf{s}_2 + \mathbf{s}_1$. The pdf for the measured data under the ith joint hypothesis may now be compactly given by

$$p(\mathbf{z} \mid H_i) = \frac{1}{\left(\pi\sigma^2\right)^N} \exp\left[-\frac{1}{\sigma^{2N}} (\mathbf{z} - S_i)^{\mathrm{H}} (\mathbf{z} - S_i) \right]. \tag{36}$$

The goal is to update the individual cell probabilities, but in general the radar cannot necessarily observe each cell apart from the others. Thus, we must first update the joint hypotheses using Bayes rule, which states that the posterior probability for each joint hypothesis is given by

$$P(H_i \mid \mathbf{z}_k) = \frac{P(H_i \mid \mathbf{z}_{k-1})\, p(\mathbf{z}_k \mid H_i)}{p(\mathbf{z}_k)} \tag{37}$$

where $P(H_i \mid \mathbf{z}_{k-1})$ is the probability of the ith joint hypothesis prior to collecting the current (kth) measurement.

The probabilities of the joint hypotheses H_i prior to the kth measurement must be calculated. Recall that there are M target parameter cells each described by a probability of a target being present in that cell. Let b_1 through b_M be the individual bits of the binary representation of a joint hypothesis. For example, the eleventh joint hypothesis in the five-cell scenario described above would have $b_5 = 0$, $b_4 = 1$, $b_3 = 0$, $b_2 = 1$, $b_1 = 1$. Since target presence or absence is assumed to be independent across cells, the probability of the ith joint hypothesis is

$$P(H_i \mid \mathbf{z}_{k-1}) = \prod_{c=1}^{M} \left(P_{c,k-1}\right)^{b_c} \left(1 - P_{c,k-1}\right)^{1-b_c}. \tag{38}$$

We desire to arrive at the updated cell probabilities. Once a measurement $\mathbf{z}_k$ is received, it is used to update the probabilities for all *joint* hypotheses. First, 2^M likelihoods must be evaluated as dictated by (36). Then all 2^M joint probabilities must be updated by (37). The updated cell

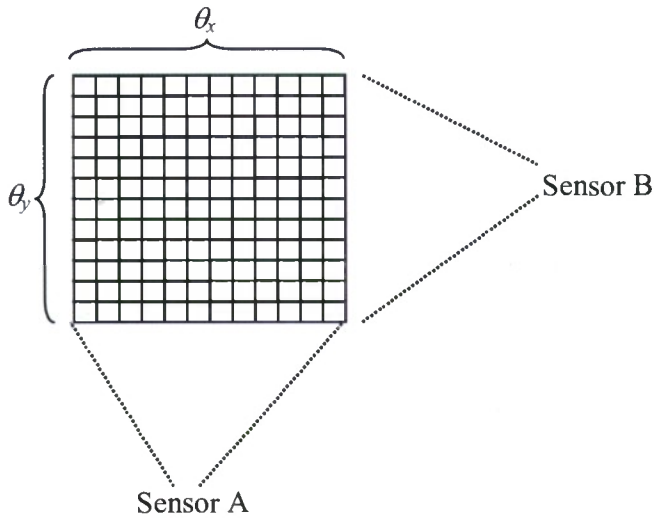

Figure 5. Two-sensor scenario with "blind" parameters.

probabilities are then obtained through the marginal probabilities of the joint hypotheses. To calculate the marginal for the *m*th cell, we sum up the probabilities for any joint hypothesis that has a target-present state for that cell. The resulting sum is the updated probability for that cell, and the process must be done for all cells.

A special case occurs when the signals produced by targets in different cells are resolvable by the radar, in which case each cell can be treated independently. In other words, while the above multiple-hypothesis testing framework applies (since it applies to all situations in general), it is computationally more efficient to perform separate probability updates for each the M cells. If the signals due to different cells are orthogonal, then separable probability updates is equivalent to the full procedure of updating joint probabilities and then calculating the individual marginal probabilities. However, the separable procedure is computationally much simpler.

We now finally consider two sensors as shown in Fig. 5. Sensor A can measure a parameter θ_x and not θ_y and vice versa for Sensor B. For example, each sensor might be used to measure angular space. As seen in Fig. 5, the channel may be described by a two-dimensional map, where each cell is described by its θ_x ,θ_y coordinates. Assuming there are M θ_x cells and M θ_y cells, there are M^2 cells in the two-dimensional parameter space. For Sensor A, note under each θ_x-cell, there are M θ_y cells that are ambiguous in the sense that Sensor A cannot resolve them. Similarly, for Sensor B, under each θ_y-cell there are M θ_x cells that are ambiguous. There are 2^{M^2} possible permutations of the overall target environment. Each permutation is a unique combination of

target presence or absence across the resolution cells, and each permutation is characterized by a probability of being true.

Because some of the cells are resolvable, we know from the above discussion that we do not have to evaluate all 2^{M^2} joint hypotheses. Rather, since each radar can resolve M cells, the hypothesis framework for each reduces into M separable updates (the resolvable cells), but embedded within each separable update are M ambiguous cells. Thus, each separable update actually requires evaluation of 2^M joint hypotheses. This is much better than evaluating 2^{M^2} joint hypotheses, but the computational issues are still formidable. Approximate, computationally reduced procedures are possible, and we have tested some, but for now the above discussions indicate a theory for updating a probabilistic grid of target cells using multiple platforms capable of resolving different regions of the parameters space.

Next, consider an example scenario of interest where Sensor A can measure or resolve parameters (θ_1, θ_2) and Sensor B measures (θ_1, θ_3). Thus, there is a common parameter (θ_1) that both sensors can measure. The SNR per measurement is 6dB and there are 4 iterations (each iteration consists of a single transmission from each of the two platforms). The initial prior probabilities are randomized on [0, 0.1] and the three-dimensional probability ensemble is formed. Because the ensemble is three-dimensional, it is difficult to plot and two of the parameters (θ_1, θ_3) are interleaved onto the same axis. In the interleaved approach, the true targets are located in cells (2,14) and (6,65) cells. Figure 6 shows the probability updates after each pair of transmissions. After four iterations, the system is confident of the presence and location of two targets. But consider what would have happened with a traditional system. The probability map after the first iteration shows confidence of target presence for the radar system after a single (in this case, a single set of two transmissions) observation. It does not appear that any targets are present, so no detections would have been declared. In fact, these two targets might never be detected by a traditional system that collects a single set of measurements, makes a hard decision, and then throws away what was learned. Only by carrying over what was learned from previous transmissions was the radar system able to finally detect the two targets. This is a fundamental benefit to sequential hypothesis testing, which is effectively what we are implementing here. Combine this carryover with optimized illumination and there exists the potential for major improvements in performance.

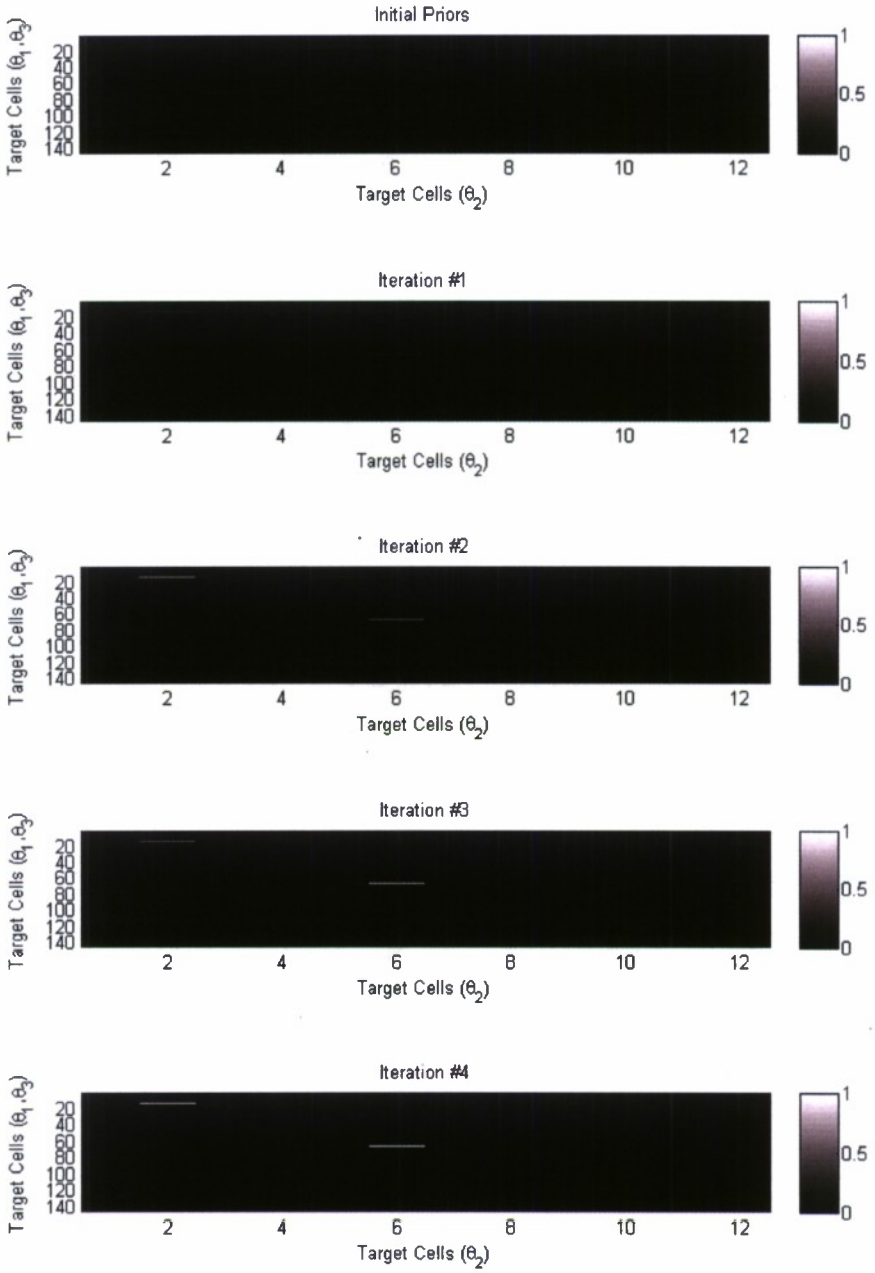

Figure 6. Three-parameter Bayesian representation with blind parameters. Probability map shown after multiple updates from two platforms.

Continuing along this theme, we computed the detection benefit that could be obtained with our strategy. Suppose that for a single data collection, a particular target has a probability of detection of 0.1 when the threshold is set to achieve a probability of false alarm equal to 0.01. Figure 7 shows a comparison of detection performance versus the number of data collections.

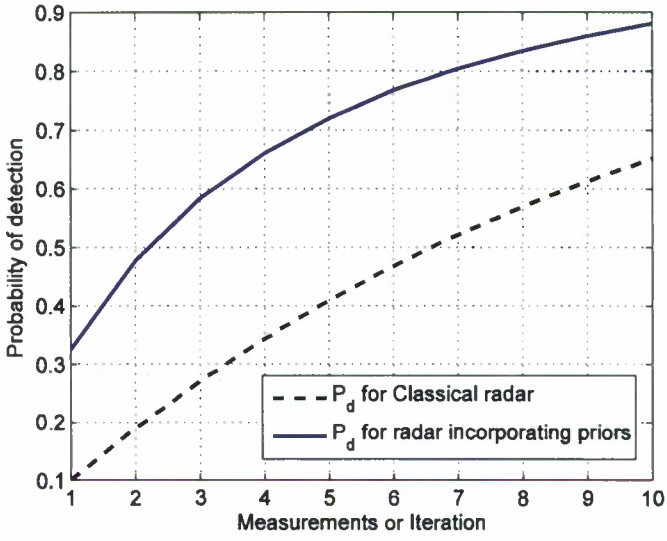

Figure 7. Detection performance of closed-loop radar vs. single-shot detection.

For the traditional system, we say the target is detected if it exceeds the threshold on at least one of the collections. For the cognitive system, we set a threshold on the probability achieved after all collections such that the probability of false alarm is the same as for the traditional system. The probability of detection is much higher for the system that carries forward what was learned on previous transmissions without even considering optimized use of time and energy through waveform control.

3. Adaptive PRF Control

A final implementation of cognitive radar that we pursued involved using our methods to optimize PRF selection. Range-Doppler ambiguities are well-known phenomena of pulsed radar systems that limit performance. These ambiguities are caused by aliasing in the time and frequency domains and their structure is controlled by pulse repetition frequency (PRF); furthermore, these ambiguities can limit effective area coverage. Solutions range from individually coding pulses to staggering the PRF, but these solutions complicate matched-filter and FFT-based range-Doppler processing.

We consider an adaptive PRF selection technique for mitigation of range and Doppler ambiguities over multiple coherent processing intervals (CPIs). Because our Bayesian framework for radar surveillance retains information obtained from previous transmissions, it is natural to consider how this information can be used to optimize PRF selection. Traditional systems either use the same PRF, or have a small set of PRFs that the system cycles through on a

regular basis. In this section, we select the PRF based on which PRF will provide the most mutual information, which in turn depends on the probabilistic channel ensemble.

Matched filtering techniques work well for range-Doppler radar when cells are unambiguous. In order to determine actual target range and Doppler for ambiguous cells, ambiguity regions are changed by adaptively selecting the PRF while target probabilities are maintained for each range-Doppler cell. In this way, we select PRFs that avoid causing uncertain cells to alias on top of each other. Applying information theory, we can determine which PRF from a pre-defined set will maximize the information obtained from the channel. Hypotheses representing all combinations of target locations within each set of ambiguous cells are formed and probabilities of each hypothesis are calculated. Due to the same computational limitations described above under the multi-platform probability update, a reduced technique is used whereby a subset of the most probable hypotheses is considered. The cell probabilities are used to calculate the variance of the ambiguous cells which is in turn used to find the conditional mutual information of the ensemble and the radar channel for given the PRF. The PRF that maximizes the mutual information (or minimizes the uncertainty) of the ensemble is chosen for the next update.

Suppose a typical constant-PRF waveform has a time-bandwidth product of *ML*. This will be the number of unambiguous range-Doppler cells, so further assume that the *ML* cells are arranged as *M* cells in one dimension (say, range) and *L* cells in the other dimension. In addition, if we desire to observe a larger area, there will be many more resolution cells that are ambiguous with the original *ML* cells. This scenario is completely analogous to the sample scenario described in the previous section where a given radar was only able to resolve target cells in one of the two target parameter dimensions. Here, we have *ML* resolvable cells, but for each resolvable cells, there are several unresolvable cells – the ambiguities. If one of the original *ML* cells shows a high probability of target presence, then that target might actually be in one of the ambiguities and be aliasing into one of the main *ML* cells. This situation will affect the probability updates and MI calculations.

Let the average RCS of a target, if it is present, be σ_s^2, and the let the maximum number of targets in a particular cell and its ambiguities be *K*. Purely speaking, the number of targets could be equal to the number of ambiguities (assuming one target per cell maximum), but limited the number of targets to *K* allows for reduced computation. Given an MxL set of orthogonal range-Doppler cells, the mutual information given a specific PRF is

$$\mathrm{MI}(\mathrm{PRF}) = \sum_{m=1}^{M}\sum_{l=1}^{L} \log_2 \pi e\left(1+\frac{\sigma_{m,l}^2}{P_n}\right) \tag{39}$$

where P_n is the noise power. The variance quantity $\sigma_{m,l}^2$ is the variance over the hypotheses of the (m,l) resolution cell. Let the average RCS of a target, if it is present, be σ_s^2. For a given cell and its ambiguities, the average combined RCS is zero for a null hypothesis (that is, no targets in the resolution cell or its ambiguous cells), σ_s^2 for a single-target hypothesis, $2\sigma_s^2$ for a two-target hypothesis, and so on up to $K\sigma_s^2$ for a K-target hypothesis. Define σ_i^2 as the combined average RCS under an i-target hypothesis. We then define

$$\sigma_{m,l}^2 = \sum_{i=0}^{K} P_{i(m,l)}\sigma_i^2 - \left|\sum_{i=0}^{K} P_{i(m,l)}\sqrt{\sigma_i^2}\right|^2. \tag{40}$$

If we become certain of how many targets are in the cell, then the corresponding P_i goes to one while the other probabilities go to zero. In this case, the variance goes to zero, and nothing more can be learned.

We use total variance in the scene to assess the performance of the MI PRF selection technique compared to other approaches. Figure 8 compares the total ensemble variance remaining in the scene after 20 CPIs for the MI-based and random-PRF selection approaches. The MI-based approach shows reduced residual variance after the 20 CPIs.

A second performance metric that can be used is the number of transmissions required to make a decision on all target parameter cells. The decision criterion is based on sequential hypothesis testing (SHT); therefore, the experiment ends when every cell in the scene satisfies

$$\frac{P_{cell}}{1-P_{cell}} > \frac{1-P_{FA}}{P_{FA}} \text{ or } \frac{1-P_{cell}}{P_{cell}} > \frac{1-P_{Miss}}{P_{Miss}} \tag{41}$$

where P_{cell} is the probability of a target being present in the cell, P_{FA} is the desired probability of false alarm, and P_{miss} is the desired probability of miss. Figure 9 compares the number of transmissions required to terminate updates for the MI and random PRF selection cases.

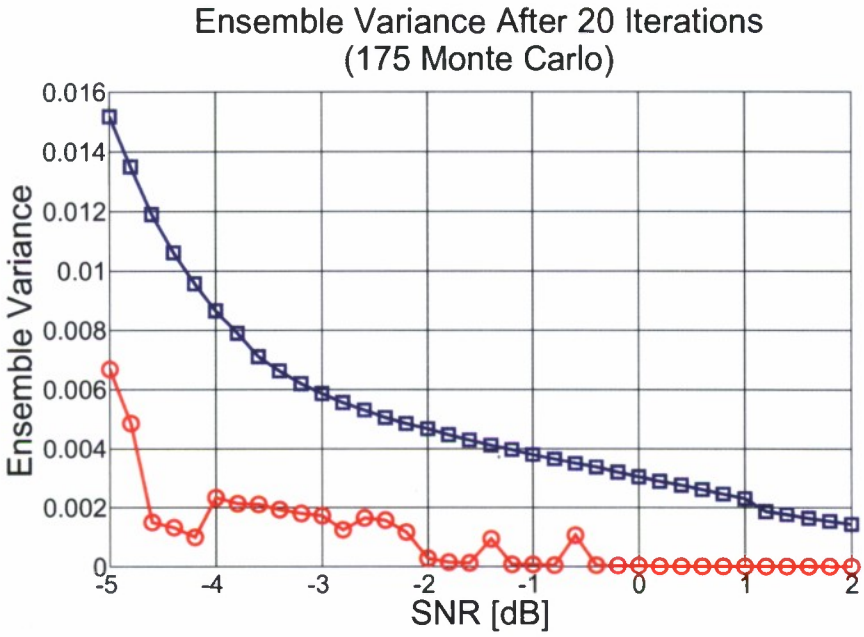

Figure 8. Comparison of variance remaining after 20 CPIs for MI-based PRF selection and random PRF selection.

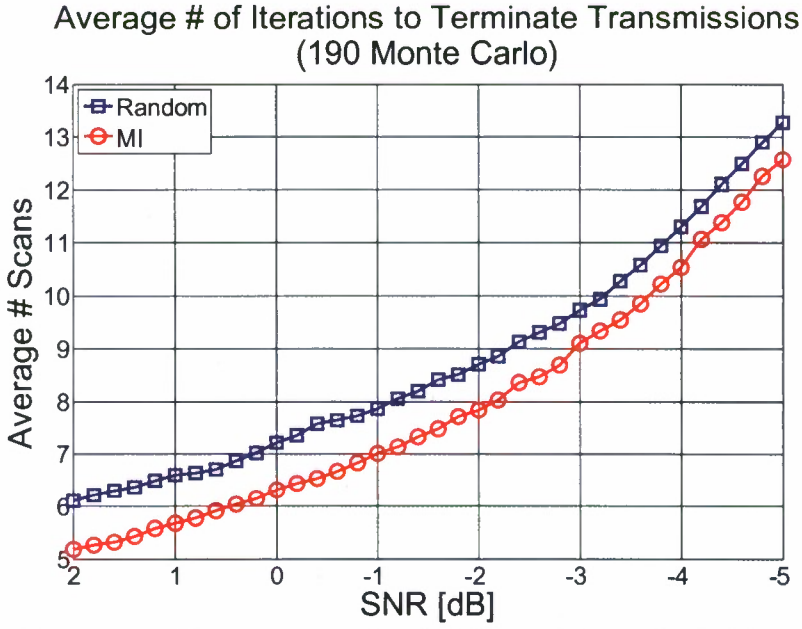

Figure 9. Comparison of the average number of CPIs required to make decisions on all cells.

Figures 10-13 show four successive snapshots of the probability ensemble for the MI-based PRF selection technique. After the first transmission, there are a number of cells that show increasecd potential for target presence when compared to the initial prior cell probability of 0.01

before any transmission are made. However, the aliasing is apparent as the probability map is periodic according to the PRF used. The PRF selection technique chooses a different PRF for the second CPI, resulting in the probability map shown in Figure 11. Many ambiguities are resolved, but there are still many uncertain cells, and we make take care not to alias them on top of each other. After the next two CPIs, two targets have been located with high confidence. In this demonstration, four PRIs were available for selection. The four PRIs were 20, 25, 30, and 40 units. To keep the CPI duration constant for each PRF, the number of pulses in the CPI was varied. Therefore, the aperture-bandwidth product of each waveform was 600, which corresponds to the number of unambiguous resolution cells (denoted by white dashed lines). The total number of cells in the scene was 8000; thus, the total cells in the scene was greater than 13x the number of unambiguous cells in any CPI, yet careful PRF selection allowed ambiguities to be resolved.

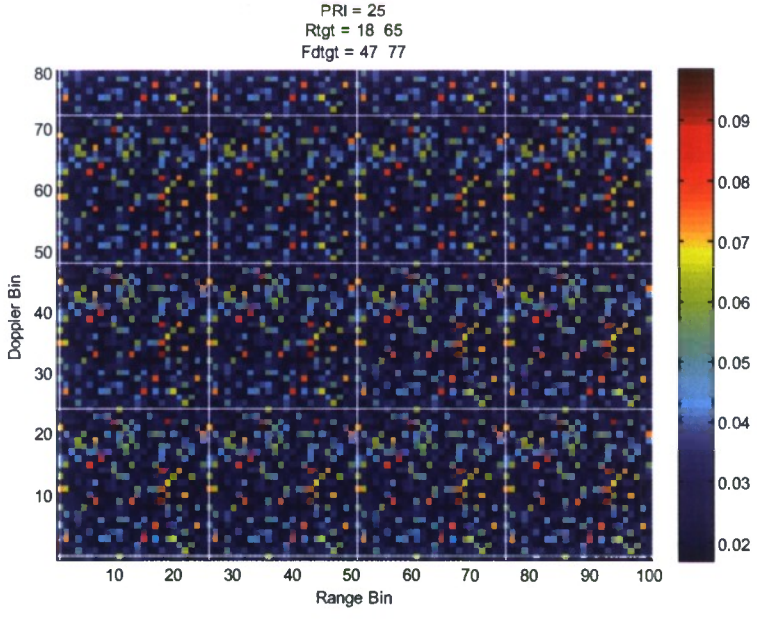

Figure 10. Wide-area probability map after first CPI.

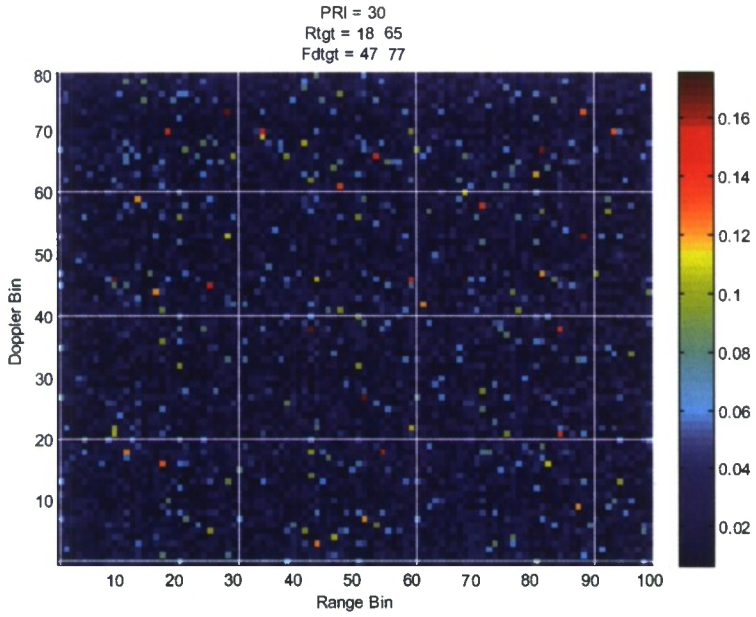

Figure 11. Wide-area probability map after second CPI.

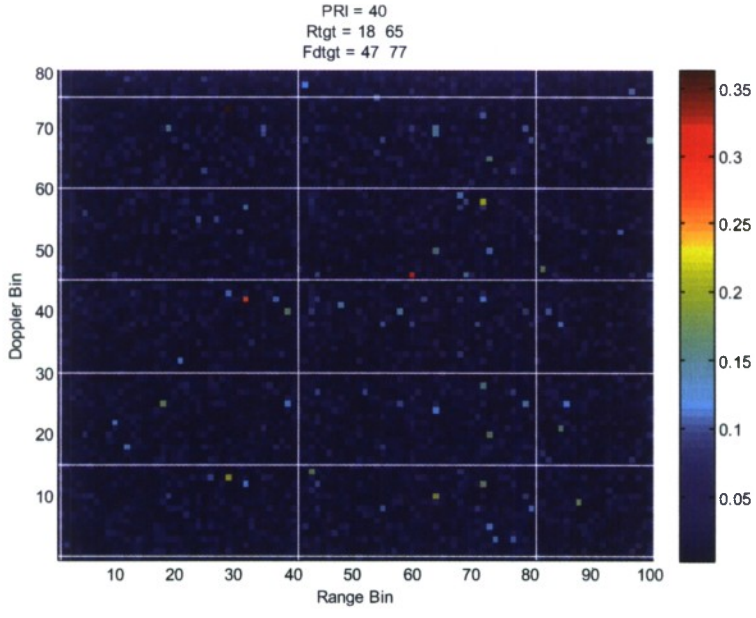

Figure 12. Wide-area probability map after third CPI.

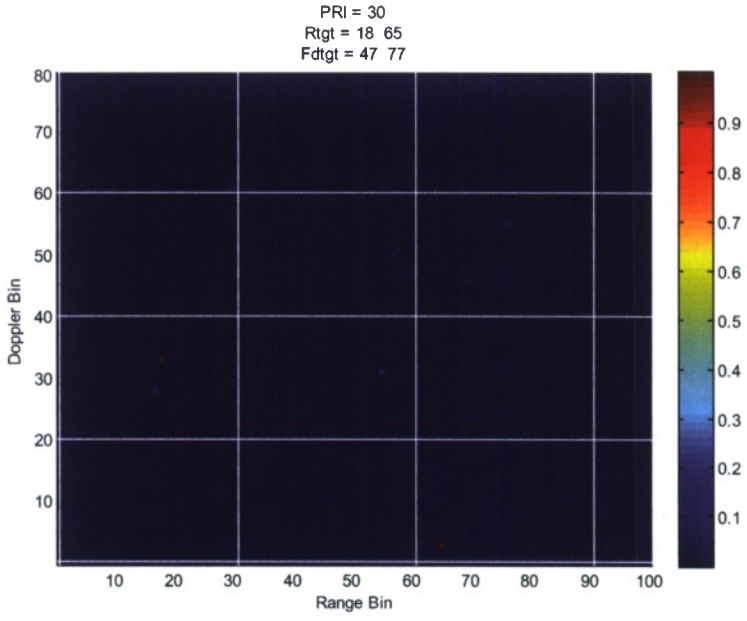

Figure 13. Wide-area probability map after fourth CPI.

4. REFERENCES

[1] M.R. Bell, "Information theory and radar waveform design," *IEEE Trans. Info. Theory*, vol. 39, no. 5, pp. 1578-1579, Sept. 1993.

[2] S.U. Pillai, H.S. Oh, D.C. Youla, and J.R. Guerci, "Optimum transmit-receiver design in the presence of signal-dependent interference and channel noise," *IEEE Trans. Info. Theory*, vol. 46, no. 2, pp. 577-584, March 2000.

[3] S. Kay, "Optimal signal design of Gaussian point targets in stationary Gaussian clutter/reverberation," *IEEE J. Sel. Topics Sig. Proc.*, vol. 1, no. 1, pp. 31-41, Jun. 2007.

[4] Ric A. Romero, Junhyeong Bae, and Nathan A. Goodman, "Theory and application of SNR and Mutual Information Matched Illumination Waveforms," accepted to *IEEE Trans. Aerospace & Electronic Systems*.

[5] S. Kay, Fundamentals of Statistical Signal Processing, Vol. I: Estimation Theory. Upper Saddle River, NJ: Prentice Hall PTR, 1993.

[6] R. Romero and N.A. Goodman, "Waveform design in signal-dependent interference and application to target recognition with multiple transmissions," *IET Radar, Sonar, and Navigation*, vol. 3, no. 4, pp. 328 – 340, August 2009.

[7] N.A. Goodman, P.R. Venkata, and M.A. Neifeld, "Adaptive waveform design and sequential hypothesis testing for target recognition with active sensors," *IEEE J. Sel. Topics in Sig. Proc*, vol. 1, no. 1, pp. 105-113, June 2007.

[8] R. Romero and N.A. Goodman, "Improved waveform design for target recognition with multiple transmissions," in *Proc. 2009 International Waveform Diversity and Design Conference*, Orlando, FL, pp. 26-30, Feb. 2009.

[9] N.A. Goodman, "Closed-loop radar with adaptively matched waveforms," in *Proc. 2007 International Conference on Electromagnetics in Advanced Applications*," Torino, Italy, pp. 468-471, Sept. 2007.

[10] J.H. Bae and N.A. Goodman, "Adaptive waveforms for target class discrimination," in *Proc. 2007 International Waveform Diversity and Design Conference*, Pisa, Italy, pp. 395-399, June 2007.

[11] Y. Yang and R. Blum, "MIMO radar waveform design based on mutual information and minimum mean-square error estimation," *IEEE Trans. Aerospace Elec. Syst.*, vol. 43, no. 1, pp. 330-343, Jan. 2007.

[12] A. Leshem, O. Naparstek, and A. Nehorai, "Information theoretic adaptive radar waveform design for multiple extended targets," *IEEE J. Sel. Topics in Sig. Proc.*, vol. 1, no. 1, pp. 42-55, June 2007.

[13] T. Butler and N.A. Goodman, "Multistatic target classification with adaptive waveforms," in *Proc. 2008 IEEE Radar Conference*, pp. 1-6, Rome, Italy, May 2008.

[14] J.H. Bae and N.A. Goodman, "Evaluation of modulus-constrained matched illumination waveforms for target identification," accepted to 2010 IEEE Radar Conference.

[15] P. Nielsen and N.A. Goodman, "Integrated detection and tracking via closed-loop radar with spatial-domain matched illumination," in *Proc. 2008 International Conference on Radar*, Adelaide, Australia, pp. 546-551, Sept. 2008.

[16] P. Nielsen, "Adaptive spatial-domain beam steering in a wide-area search and track application," M.S. Thesis, ECE Dept., University of Arizona, 2009.

5. PUBLICATIONS RESULTING FROM GRANT

Accepted:

R. Romero, J.H. Bae, and N.A. Goodman, "Theory and application of SNR- and MI-based matched illumination waveforms," *IEEE Trans. on Aerospace and Electronic Systems*, to appear.

R. Romero and N.A. Goodman, "Waveform design in signal-dependent interference and application to target recognition with multiple transmissions," *IET Radar, Sonar, and Navigation*, vol. 3, no. 4, pp. 328 – 340, August 2009.

N.A. Goodman, P.R. Venkata, and M.A. Neifeld, "Adaptive waveform design and sequential hypothesis testing for target recognition with active sensors," *IEEE J. Selected Topics in Signal Processing*, vol. 1, no. 1, pp. 105-113, June, 2007.

R. Romero and N.A. Goodman, "Improved waveform design for target recognition with multiple transmissions," in Proc. 2009 International Waveform Diversity and Design Conference, Orlando, FL, pp. 26-30, Feb. 2009.

P. Nielsen and N.A. Goodman, "Integrated detection and tracking via closed-loop radar with spatial-domain matched illumination," in Proc. 2008 International Conference on Radar, Adelaide, Australia, pp. 546-551, Sept. 2008.

T. Butler and N.A. Goodman, "Multistatic target classification with adaptive waveforms," in Proc. 2008 IEEE Radar Conference, pp. 1-6, Rome, Italy, May 2008.

R. Romero and N.A. Goodman, "Information-theoretic matched waveform in signal-dependent interference," in Proc. 2008 IEEE Radar Conference, pp. 1-6, Rome, Italy, May 2008.

N.A. Goodman, "Closed-loop radar with adaptively matched waveforms," in Proc. 2007 International Conference on Electromagnetics in Advanced Applications," Torino, Italy, pp. 468-471, Sept. 2007.

J.H. Bae and N.A. Goodman, "Adaptive waveforms for target class discrimination," in Proc. 2007 International Waveform Diversity and Design Conference, Pisa, Italy, pp. 395-399, June 2007.

Submitted:

R.A. Romero, C.M. Kenyon, and N.A. Goodman, "Channel probability ensemble update for multiplatform radar systems," submitted to 2010 Cognitive Information Processing (CIP).

J.H. Bae and N.A. Goodman, "Evaluation of modulus-constrained matched illumination waveforms for target identification," accepted to 2010 IEEE Radar Conference.

H.S. Kim and N.A. Goodman, "Waveform design by task-specific information," accepted to 2010 IEEE Radar Conference.

Book Chapters:

N.A. Goodman, P. Venkata, and R. Romero, "Iterative Technique for System Identification with Adaptive Signal Design," to appear in Applications and Methods of Waveform Diversity, SciTech Publishing.

N.A. Goodman, J.H. Bae, and R. Romero, "Waveform Design for Target Class Discrimination with Closed-Loop Radar," to appear in Applications and Methods of Waveform Diversity, SciTech Publishing.

6. Inventions and Patent Disclosures

None.